Router Hardware

An important requirement for establishing any peering link is to ensure that there is suitable router hardware available. This section looks at what considerations need to be made when choosing a router for the new connection being planned.

For completeness, the section will also cover the router consideration for the entity's first Internet connection (to its upstream provider).

- First Internet Link (Transit)
- Private Peering Link
- Public Peering Link

First Internet Link

For an organisation embarking on establishing their first Internet connection (to their transit provider), a router will be required. The Peering Toolbox can't provide an exhaustive summary of all the options and combinations available as they depend on circumstances, local conditions and market, but the key points to note are documented here.

- Router Type
- Interface Considerations
- Router Throughput
- IPv4 & IPv6
- BGP needs
- Packet Filtering

Type

The Peering Toolbox is aimed at organisations who are or are planning to take part in peering. For this reason a consumer/home router (often erroneously called an “internet modem”) is not sufficient and cannot be recommended (even through there are some quite capable devices available).

The type that needs to be looked at need to be “enterprise grade” which means the router is offered with reasonable warranty, a support contract (often required by enterprises), is usefully rack (or shelf) mountable, has sufficient cooling, has possibility of redundant power supplies, has a console port for in-situ access, is accessible remotely using Secure Shell, supports SNMP (Simple Network Management Protocol), and supports a command line interface suitable for human or automated tool use.

Alternatively, there are several software routers available which could be installed on a Linux container or virtual machine or small Linux appliance. These might be entirely suitable, as the software is usually fully featured (like the main stream vendor routers) and very capable. One example of a software router is FRR which is widely used as an alternative to dedicated vendor supplied hardware.

Interfaces

The router needs to have internal and external interfaces to match the physical media in use.

Internal interface types are usually only Gigabit Ethernet today, at a minimum, even on the most inexpensive devices. The number of internal interfaces needed depends entirely on the organisational needs. Usually a single or dual interfaces are all that are needed, connecting to the organisation's core router, or routers. The diagram shows two possibilities - the core router drawn represents the existing network infrastructure, be it an existing router or layer-3 switch.

It is recommended to separate the router that has the upstream connection from the core of the network, for the security of having a clear demarcation point, and so that specific border functions don't overload the core devices in the network.

External interface types can range from Gigabit Ethernet, SFP-based fibre optics, various coaxial or copper telephony wiring, to point-to-point wireless. Many of the enterprise routers come with a dizzying amount of configuration options - if the future trajectory for the upstream (and peering) physical media access is uncertain, perhaps specifying a router that has a range of upstream link options would be the most prudent choice to make.

Throughput

The router needs to be able to handle the throughput of the link being purchased, and leave sufficient CPU and memory capacity for future upgrades.

Note that even if the router may have Gigabit or fibre optic interfaces, there is no guarantee that it can actually deliver Gbps rates. This is especially true for the CPU based routers whose throughput slows down significantly with increasing traffic, the amount of packet filtering configured, and Network Address Translation (if sufficient IPv4 address space is not available).

It is important to check with the vendor what the true throughput is in a realistic use case (known as *Internet Mix* or IMIX, representing the typical average packet size seen on the Internet today), not lab testing!

IPv4 & IPv6

If the transit provider has deployed both IPv4 and IPv6 on their network and offers the capability to customers, it is strongly recommended that the router chosen be able to handle IPv4 and IPv6, known as dual-stack operation. This dual-stack support needs to have all commands available supporting both IP protocols (and not a reduced command set for IPv6).

Using IPv6 is advantageous as it means that content traffic (which forms about 80% of typical Internet traffic today) will not have to traverse Network Address Translation devices in the upstream's network or use the NAT feature on the router, reducing the resource burden, and also improving the service quality experienced by the end-users.

BGP

(UPDATED)

Most “first time” Internet connections will simply use a static default route pointing to the upstream provider, with the upstream pointing a route to their customer for the customer's address space.

However, it pays to think forwards, especially considering that this Toolbox is all about how an organisation should go about peering! And for that, BGP will be required, and it is recommended that any new procured router is fully BGP capable.

Some end-sites will start off with using BGP even for their first Internet connection, from day one. Historically they'd use a private AS number for this, but with the relaxation of policies in some of the Regional Internet Registry regions, a public AS number can now be obtained simply by becoming a member of the RIR and receiving address space.

If BGP is going to be used on the link, the router must be BGP capable, although it does not have to or need to carry the full BGP table (which is large and growing rapidly). Most modern routers have implemented the latest BGP standards and extended capabilities - reviewing BGP Best Practices documentation and comparing with vendors' claimed feature support is strongly recommended.

If BGP is being used on this transit link, and there are no other external links for this network, then all the operator needs to do is announce their address space to their upstream, and accept a default route from their upstream. This scenario is discussed in the Single Upstream section of the Toolbox.

The control plane needs of the router (the processor to handle BGP and other routing protocols) are not significant where there are just a few BGP peers and only a few routes are being handled as in this case.

Packet Filtering

The final router requirement is the ability to do packet filtering, with at least the ability to filter by source address, destination address, source port, destination port, and IP protocol.

It is important to check how many of these filter rules the router will support, and if performance degrades as more rules are added. Ideally there should be minimal performance impact as rules are added; be aware that CPU based routers are likely to show a significant performance hit as rules are added.

The minimum filtering needed on an enterprise connection today would be:

- allow all ICMP
- allow inbound established TCP connections (sessions originated internally)
- allow externally originated connections inbound to public hosted services (website, email server)
- block external access to network infrastructure control planes
- allow outbound traffic only from public address space used internally (anti-spoofing)
- allow UDP such that essential UDP based services work (Domain Name Service, Network Time Protocol, etc)

A network operator will likely be more generous, with filter rules only blocking access to the network infrastructure control planes, implementing anti-spoofing filters, but permitting all other public address space.

A detailed discussion of filter rules is beyond the scope of the Peering Toolbox.

Private Peering Link

When implementing the first private peering link, it is recommended to procure a separate router for this function. This router is normally dedicated only for peering connections, whether connecting to public peers at an IXP or private peers.

If procuring a separate router is not a possibility, it is possible that an existing router could be used, so long as it meets the appropriate technical requirements for participating in a peering infrastructure (full support of BGP, sufficient control-plane memory and CPU capacity). There is a security caveat with using one router for connecting to both a transit provider and a peer though - that router will have a default route which could potentially be abused by the private peer (who could simply point a default route at it, and get “free” outbound transit - if they wanted to).

This section gives recommendations on the assumption that the organisation will be procuring a separate router for their private peering connection.

Type

The discussion about the type of router used in the Transit Connection applies here too. There are no “lesser requirements” of class of router when implementing a peering connection.

Interfaces

The interface consideration is the same as for the router being used for the Transit Connection. As a separate router is being procured, the connectivity diagram might end up looking like those shown below.

Throughput

The discussion about the throughput of the router used in the Transit Connection applies here too. In fact a peering router, longer term, will likely have higher normal throughput requirements than a

transit router, especially if the operator reaches their goal of achieving 80% of their traffic by peering. Of course, in the event of the peering connectivity failing, the transit router has to be able to carry the load (up to the capacity of the transit link).

IPv4 & IPv6

The discussion about the type of router used in the Transit Connection applies here too. If the operator has deployed IPv6 in addition to IPv4 to their upstream provider, then naturally the router procured for the peering link needs full dual-stack support as well.

When operating a dual stack network, it is strongly recommended to ensure that whatever connectivity is supplied for IPv4 is also replicated for IPv6. For example, if the peering links are IPv4-only, yet the transit is dual stack IPv4/IPv6, then potential peering traffic will use the paid-for transit link, rather than the free peering link.

BGP

(UPDATED)

BGP will be required for any peering connection, in which case the peering router has to fully support BGP.

The main difference between the needs for a peering link and the Transit Connection discussed earlier is that the peering router will only ever carry the peer routes and internal routes of the organisation. In fullness of time, especially for network operators, and large enterprises, the transit router could quite likely end up carrying a large portion if not the entire BGP table, so will have to be sized accordingly. The peering router, on the other hand, can have a more modest memory (both RAM for the routing table and forwarding table) requirement.

The BGP configuration used on a private peering connection is discussed in the Single Upstream and Private Peer section of the Toolbox.

The control plane needs of the router (the processor to handle BGP and other routing protocols) are not significant where there are just a few BGP peers and only a few routes are being handled as in this case.

Packet Filtering

The discussion about Packet Filtering support for the router used in the Transit Connection fully applies here too.

Public Peering Link

This router is normally dedicated only for peering connections, whether connecting to public peers at an IXP or private peers. If procuring a separate router is not a possibility, it is possible that an existing router could be used, so long as it meets the appropriate technical requirements for participating in a

peering infrastructure (full support of BGP, sufficient controlplane memory and CPU capacity).

Type

The discussion about the type of router used in the Private Peering Link applies here too. There are no “lesser requirements” of class of router when implementing a public peering connection.

In fact, for a public peering connection, there is a possibility that the router will have to be installed at the Internet Exchange Point itself, which might have more stringent requirements about cooling, being rack mountable, remote management access, physical console access, and power.

Interfaces

The interface consideration is the same as for the router being used for the Private Peering Link. As a separate router is being procured, the connectivity diagram might end up looking like either of those shown below:

In more sophisticated cases, especially where the IXP is providing a significant portion of the organisation's traffic, a second peering router might be procured, to provide not only link redundancy but also router redundancy. This second/redundant infrastructure will likely only be deployed a bit further down the track, after gaining sufficient operational experience of traffic levels and infrastructure reliability. The digram below shows this scenario:

In the case the operator has private peers as well, these could connect to the peering router at the IXP (the Private Network Interconnect or PNI described elsewhere). Or, the operator would have a separate on-premises peering router.

Throughput

The discussion about the throughput of the router used in the Private Peering Link applies here too.

IPv4 & IPv6

The discussion about IPv4 and IPv6 support for the router used in the Private Peering Link fully applies here too.

BGP

(UPDATED)

The discussion about BGP support for the router used in the Private Peering Link fully applies here too.

The BGP configuration used on a public peering connection is discussed in the Single Upstream and IXP section of the Toolbox.

The control plane needs of the router (the processor to handle BGP and other routing protocols) in this case can be quite significant and care is needed when selecting suitable hardware.

Small IXPs will have only a few peers so there it a likelihood that only a few thousand routes will be received by the new member. Most standard router hardware has sufficient control plane capability to handle this.

Larger IXPs will likely have dozens of members, with the largest IXPs today approaching one thousand members. This has significant control plane demands on the peering router, and it is important that one with a powerful control plane CPU is chosen, especially one that is proven to handle several hundred peers and tens of thousands of routes with ease. In this case it can be helpful to consult with existing IXP members seeking suggestions or recommendations. Note the usual caveat between vendor marketing claims and real world experience of network operators.

Packet Filtering

The discussion about Packet Filtering support for the router used in the Private Peering Link fully applies here too.

Back to "What is required for Peering" page